

UNITED STATES PATENT APPLICATION

FOR

ERROR DIFFUSION OF MULTIREOLUTIONAL REPRESENTATIONS

Inventor:

Maya Gupta

Prepared by:

Blakely, Sokoloff, Taylor & Zafman
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025
(408) 720-8300

Attorney's Docket No. 074451.P135

EXPRESS MAIL CERTIFICATE OF MAILING


"Express Mail" mailing label number: EL867648006US

Date of Deposit: December 19, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service
"Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has
been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Melanie Lyons

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

12-19-01
(Date signed)

ERROR DIFFUSION OF MULTIREOLUTIONAL REPRESENTATIONS

FIELD OF THE INVENTION

[0001] This invention relates generally to the field of image processing, and more specifically to perform error diffusion on multiresolutional representations of images.

BACKGROUND OF THE INVENTION

[0002] Error diffusion is an idea generally credited to Floyd and Steinberg (See "An adaptive algorithm for spatial greyscale," R.W. Floyd and L. Steinberg, Proceedings of the Society for Information Display, vol. 17, no.2, pp. 75-77, 1976) and dates back to 1976, though it is equivalently a two-dimensional implementation of sigma-delta modulation, which as a tool for one-dimensional A/D converters dates at least to 1963 (see "A unity bit coding method by negative feedback," H. Inose and Y. Yasuda, Proceedings of the IEEE, vol. 51, pp. 1524-1535, Nov. 1963). As described therein, the basic error diffusion process is to raster scan the pixels, and for each pixel quantize it to black or white, then pass the error to neighboring pixels. However, the process is done serially, and for each pixel, the error between the original pixel and the quantized version is calculated and spatially filtered to pixels that have yet to be quantized. Thus pixels that are quantized are not the original pixels from the image, but the original pixels plus the error they have inherited from pixels previously quantized. Figure 8 is a flow diagram showing a visual representation of the error diffusion process. Note that error diffusion is also equivalent to DPCM Differential Pulse Coded Modulation.

[0003] Error diffusion performs fairly well and is used often in practice. It is particularly good at accentuating edges in an image. Possibly the worst side effect is the graininess of a resulting image. In most applications, error diffusion quantizes between white and black, but error diffusion allows for a more general quantization. For instance, vector error diffusion of color images quantizes each incoming pixel to the closest of a small number of colors in a preset palette, where each color is defined by a three-dimensional vector.

[0004] Researchers have manipulated many factors in error diffusion, including the filter taps, the raster scan order, spatially-varying filters, adaptive filters approaches mixed with clustered dot dither, dot gain compensation, etc.

[0005] Halftoning takes a continuous image with many colors and creates an image that looks like the original but uses only two or a few colors. Halftoning is often implemented using the error diffusion algorithm, which is performed sequentially on pixels. To date, error diffusion is an algorithm performed on pixels, in the pixel domain.

[0006] Two papers, "Adaptive Error Diffusion for Multiresolution Rendering," P.W. Wong. Society for Information Display Digest, 1994, and "Embedded Multilevel Error Diffusion," J.R. Goldschneider, E. Riskin, P.W. Wong. IEEE Transactions on Image Processing, vol. 6, no. 7, July 1997, investigate the problem of designing an error diffused image that can be well rendered at many different resolutions. However, all of the operations therein are done in the pixel domain.

[0007] Other sources, including "Wavelet methods for compression, rendering, and descreening in digital halftoning," N. A. Breux and C. H. Chu. SPIE Conference Proceedings, vol. 3078, 1997, and "Multiresolution error-convergence halftone algorithm," E. Peli. Journal of the Optical Society of America A, vol. 8, no. 4, April 1991, attempt to match the energy of each level in the halftone to the energy in that level of the multiresolutional original image. The goal of these sources is apparently to yield a halftone that is designed for progressive display at increasing resolution. These halftoning techniques do not use error diffusion.

[0008] One conference paper, "A Recursive Multiscale Error-Diffusion Technique for Digital Halftoning," I. Katsavoundidis and C.C. Kuo, SPIE Conference Proceedings vol. 2094, 1993, also proposes the goal of multi-level mean-matching. This paper describes pixels being quantized, but in a random scan order. Error is diffused symmetrically about a quantized pixel, as opposed to the normal forward-only error propagation approach. However, because some neighboring pixels may already be quantized, the error that would have been passed to them is summed into a "global error" for that local region. These "global errors" are passed to neighboring regions using the same spatial error filter. The paper proposes the idea of passing some error from different regions within a scale. Again, however, it is the pixels themselves which are quantized.

[0009] The same authors refined their ideas and presented a modified algorithm in "A multiscale error diffusion technique for digital halftoning," I. Katsavoundidis and C.C. Kuo, IEEE Transactions on Image Processing, vol. 6, no. 3, March 1997. That algorithm employs a nonlinear but non-random scanning order of the pixels, but still

quantizes the pixels themselves, and still uses a non-causal filter (which is made possible due to the nonlinear scan order). The multiresolutional framework only affects the scan order of the pixels.

[0010] In "Neural Network Adaptive Digital Image Screen Halftoning (DISH) based on Wavelet Transform Preprocessing," H. Szu, Y. Zhang, M. Sun, and C.C. Li, SPIE Conference Proceedings, vol. 2242, 1994, an adaptive screen for dithering was proposed in which the information about how to adapt the screen is pulled from the wavelet transform representation.

[0011] Speed of copying, printing, and processing are especially important to business customers. More image processing is moving to wavelet, or multiresolutional domains. Yet the image has traditionally had to be converted back to the pixel domain to do the error diffusion. Speed could be increased if this conversion could be avoided.

SUMMARY OF THE INVENTION

[0012] A method and apparatus for performing error diffusion of a multiresolutional representation of an image is disclosed herein. In one embodiment, an image is received as a pixel representation. The image is converted from a pixel representation to a multiresolutional representation having multiresolutional coefficients. Each multiresolutional coefficient is quantized and an error is passed to other multiresolutional coefficients as specified by an error filter. Finally, a resulting quantized image is produced as an output.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0014] **Figure 1** is a flow diagram of one embodiment of a process for performing error diffusion of a multiresolutional representation of an image;

[0015] **Figure 2** is a flow diagram of one embodiment of a process for performing error diffusion that includes arranging the multiresolutional coefficients to represent a spatial dependence of each multiresolutional coefficient for error diffusion of a multiresolutional representation of an image and scaling the multiresolutional coefficients to contain a total energy of a region represented by each multiresolutional coefficient for error diffusion of a multiresolutional representation of an image;

[0016] **Figure 3** is a flow diagram of one embodiment of a process for performing error diffusion of a multiresolutional representation of an image in which an error filter is adapted to a multiresolutional level represented by a multiresolutional coefficient creating the error;

[0017] **Figure 4** is a flow diagram of one embodiment of a process for performing error diffusion that includes passing the error generated using an error filter that depends on the subband of the coefficient whose error is being passed;

[0018] **Figure 5** is a flow diagram of one embodiment of a process for performing error diffusion that includes passing the error generated using an error filter that

depends both on the subband and on the multiresolutional level of the coefficient whose error is being passed;

[0019] **Figures 6A and 6B** are block diagrams of embodiments of a device to perform error diffusion of a multiresolutional representation of an image;

[0020] **Figures 7A and 7B** are diagrams comparing an example of the output of an error diffusion filter applied to the pixel domain with an example of the output of an error diffusion filter applied to the wavelet domain;

[0021] **Figure 8** is a block diagram showing a visual representation of error diffusion;

[0022] **Figure 9A** illustrates an arrangement of first-level Haar coefficients for a one-level Haar wavelet transform;

[0023] **Figure 9B** illustrates an arrangement of two-level Haar coefficients for a two-level Haar wavelet transform

[0024] **Figures 10A and 10B** illustrate an example of sealing the difference coefficients additionally to provide extra sharpening; and

[0025] **Figure 11** illustrates one embodiment of a wavelet coefficient with an indication of wavelet coefficients that are spatially close.

DETAILED DESCRIPTION

[0026] The present invention performs error diffusion on multiresolutional representations of images. In one embodiment, the error diffusion is performed in the wavelet domain. In this case, during the error diffusion process, wavelet coefficients in the multiresolutional representation are quantized and one or more error values representing, at least in part, a difference between the quantized coefficients and non-quantized coefficients, are passed onto one or more other (e.g., spatially neighboring) coefficients.

[0027] By removing the need to invert the wavelet transform before halftoning, increased copying and printing speed can be achieved. Parallel error diffusion allows for a huge increase in speed, and parallel error diffusion may be a more realistic technique if done on wavelet coefficients due to the non-causality of error diffusion on the wavelet domain which crosses the artificial borders created by parallel error diffusion. Apart from increases in processing speed, it may be possible to design new wavelet domain error propagation spatial filters that achieve sharp edges and fewer "wormies" (artifacts) than is now possible. Wormies are an artifact where halftone dots bunch together in long squiggly lines, and they are a common problem in error diffusion halftones.

[0028] In the following description, numerous details are set forth to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific

details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0029] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0030] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system

memories or registers or other such information storage, transmission or display devices.

[0031] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0032] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0033] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For

example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

[0034] Figure 1 is a flow diagram showing a method for error diffusion of a multiresolutional representation of an image. The process is performed by processing logic that may comprise hardware (e.g., circuitry, dedicated logic, etc.), software (such as is run, for example, on a general purpose computer system of a dedicated machine), or a combination of both.

[0035] Referring to Figure 1, in processing block 101, processing logic receives an image as a pixel representation. In processing block 102, processing logic converts the image from a pixel representation to a multiresolutional representation having multiresolutional coefficients. The multiresolutional representation may be, in varying embodiments, a pyramidal decomposition, a wavelet decomposition, a time-frequency decomposition, or another multiresolutional representation of an image. In one embodiment, the decompositions may be two-dimensional (2-D). In an alternative embodiment, the decompositions may be three-dimensional (3-D).

[0036] Multiresolutional coefficients may include scaling coefficients and wavelet (difference) coefficients. These may be referred to herein as transform coefficients. In processing block 103, processing logic quantizes each multiresolutional coefficient and passes one or more error values to spatially close multiresolutional coefficients as specified by an error filter. In other words, processing logic adds that one or more error

values to spatially neighboring coefficients. In this manner, error diffusion is performed on the multiresolutional coefficients.

[0037] In essence, wavelet domain coefficients contain information about local averages and differences, or "errors to the average." Diffusing the error spreads out the local averages. The difference coefficients cause the image's edge information to appear in the halftone.

[0038] Any quantizer may be used to perform the quantization. For example, a simple quantizer may be used which performs threshold on each coefficient and outputs one value or another based on whether the coefficient is above or below a threshold respectively. For instance, the quantizer may be as follows for each coefficient x_{input}

if $x_{\text{input}} < 0.5$, then $x_{\text{output}} = 0$;
otherwise, $x_{\text{output}} = 1$.

Such a quantizer is particularly useful for bilevel outputs (like Black-White halftoning). In another embodiment, a multi-level quantizer may be used, such as, for example, the following, where $y_i \in (y_1, y_2, \dots, y_l)$, the multiple levels of the quantizer.

$x_{\text{output}} = y_i$ that minimizes $(y_i - x_{\text{input}})^2$.

For more information on this quantizer, see A. Gersho & R. Gray, Vector Quantization and Signal Processing, Academic Publishers 1992, Massachusetts.

[0039] The errors produced as a result of quantization are passed onto neighboring coefficients. In one embodiment, the errors are passed according to the following table:

	X	7/16
3/16	5/16	1/16

That is, the error incurred from coefficient X is diffused forward onto unprocessed coefficients on the right and below according to the weights given above.

[0040] In another embodiment, the spatially close multiresolutional coefficients to which the error is passed may be adjacent coefficients, spatially nearby coefficients, or coefficients away from the coefficient being quantized. Figure 11 illustrates one example of a wavelet coefficient X which represents the coefficient that has just been quantized and other coefficients that are to be considered spatially close indicated with cross-hatching. Thus, the error may be passed to adjacent coefficients in the same or neighboring rows as well as coefficients in nearby rows of which the coefficients themselves are not neighboring coefficients nor are they in neighboring rows.

[0041] Thus, in general, let e_{ij} , an error produced at some (i,j) spatial location, be

$$e_{ij} = x_{ij} - \hat{x}_{ij}$$

where x_{ij} equals the original multiresolutional coefficient and \hat{x}_{ij} equals a quantized version of x .

[0042] In one embodiment, the error filtering operation then changes the coefficients to include the error:

$$\begin{aligned}
x_{i,j+1} &= x_{i,j+1} + \frac{11}{27} e_{ij} \\
x_{i,j+2} &= x_{i,j+2} + \frac{1}{27} e_{ij} \\
x_{i+1,j} &= x_{i+1,j} + \frac{11}{27} e_{ij} \\
x_{i+1,j+1} &= x_{i+1,j+1} + \frac{3}{27} e_{ij} \\
x_{i+2,j} &= x_{i+2,j} + \frac{1}{27} e_{ij}
\end{aligned}$$

[0043] Note that some errors may be positive, some negative, and some transform coefficient may be negative.

[0044] Finally, in processing block 104, processing logic produces a resulting quantized image as an output. By performing error diffusion in the wavelet domain, an output similar to error diffusion in the pixel domain is created, but it is theoretically different, and in practice the results are also different. Theoretically, a key difference is that in the process of error diffusion, when a pixel is quantized it contains information only about the pixels that came before it (because it has received the errors of these predecessors). However, when error diffusion is performed in the wavelet domain, each coefficient already represents information about the pixels around it, including pixels that have not yet been quantized. A main advantage of this approach is the reduced processing time required to produce an output due to the removal of the need to return to the pixel domain before performing error diffusion.

[0045] Figures 7A and 7B provide an example the output of an error diffusion filter applied to the pixel domain with an example of the output of an error diffusion filter applied to the wavelet domain. Image 701 is the output of an error diffusion filter

applied to a pixel domain. Image 702 is the output of the same error diffusion filter applied to a wavelet domain; in this example, a Haar wavelet domain. It is shown that the same filter creates different outputs depending on whether the filter is applied to the wavelet or pixel domain. "Wormies," artifacts that look like worms in smooth regions, are different in the outputs. Furthermore, there is less hemming (i.e., false white surrounding dark text) in the Haar domain image 702. This makes the Haar domain image 702 look less jarring than the pixel domain image 701. Performing the error diffusion on the space-frequency (time-frequency) representation may provide similar image quality and other advantages, such as reduced computation time and faster output, and smoother background regions.

[0046] Figure 2 is a flow diagram of one embodiment of a process for performing error diffusion that includes arranging the multiresolutional coefficients to represent a spatial dependence of each multiresolutional coefficient for error diffusion of a multiresolutional representation of an image. In this embodiment, converting the image from a pixel representation to a multiresolutional representation having multiresolutional coefficients, shown in process block 102, includes arranging the multiresolutional coefficients to represent a spatial dependence of each multiresolutional coefficient, shown in included process block 201.

[0047] Figure 9A illustrates an arrangement of first-level Haar coefficients for a one-level Haar wavelet transform.

[0048] Figure 9B illustrates an arrangement of two-level Haar coefficients for a two-level Haar wavelet transform.

[0049] Note that arranging the pixels allows the image to look correct, and not appear disjointed and distorted. The arrangement is performed to get the transform coefficient “close” to the spatial area about which it has information about, otherwise, your image would come out disjointed and distorted.

[0050] Creating a halftone as described herein is different than error diffusion on the pixel domain. The same effect cannot be captured by error diffusing on the pixel domain. One difference is the causality, or how the information is being propagated. Error diffusion on the pixel domain only propagates information forward (in space); that is, it is a causal algorithm. The above embodiment first transforms the image from a pixel domain to a multiresolutional domain, which means that some coefficients now contain information about pixels in the future (spatially). For instance, a LL coefficient from a one level Haar transform, as shown in Figure 9, contains information about pixels in the region corresponding to the LL coefficient. Thus, future information about the pixels is accessed when it quantizes that LL coefficient. Thus, the technique described herein is non-causal with respect to the pixel domain. It follows that no error diffusion done directly on the pixel domain could create the same result due to the causal nature of such direct error diffusion.

[0051] Viewed in light of causality, the above embodiment can be correctly said to preprocess the pixel image (i.e., by taking local averages and differences) and then error diffuse it.

[0052] Referring back to Figure 2, the process of converting the image from a pixel representation to a multiresolutional representation having multiresolutional

coefficients, shown in process block 102, includes scaling the multiresolutional coefficients to contain the total energy of a region represented by each multiresolutional coefficient, shown in included process block 202. The scaling may be done to ensure that smooth regions don't end up having the wrong average color. Such scaling may be performed as described in U.S. Patent Application Serial No. 09/467,544, entitled "Multiscale Sharpening And Smoothing With Wavelets", filed December 10, 1999, which is incorporated herein by reference and assigned to the corporate assignee of the present invention. One could choose to scale the difference coefficients by more, and this provides extra sharpening.

[0053] Figures 10A and 10B illustrate an example of scaling the difference coefficients additionally to provide extra sharpening. Referring to Figure 10A, image 1001 is the result of performing wavelet error diffusion as described above without scaling. Image 1002 in Figure 10B illustrates the result of performing wavelet error diffusion with wavelet difference coefficients scaled up by 3.

[0054] Thus, a halftone is created from a multiresolutional representation of the image using error propagation and without inverting the multiresolutional transform. The result is that the output quantized image may be a good halftone of the original pixel image. An intuitive way to understand the process is by viewing the multiresolutional representation as a downsampled image (i.e., the lowest resolution) with errors (i.e., the difference or wavelet coefficients). Each coefficient is quantized, and then passes the coefficient's influence (error) on, so that the lowpass coefficients influence an entire region by dint of the error propagation. The difference coefficients

attempt to locally correct that influence to ensure the original image's edges are visible in the halftone, and if the correction is not complete, then there will be an error that is passed on to correct it nearby.

[0055] In one embodiment, a one level Haar transform is used to create the multiresolutional representation and binary quantization is subsequently applied thereto. This results in the generation of an output that is a good halftone of the original pixel image. In alternative embodiments, the number of multiresolutional levels increases. As the number of multiresolutional levels increases, more artifacts may be introduced due to the multiresolutional representation, and a more distorted looking halftone may result. One possible solution is to use multi-level error filters and vary the error spatial filter depending on what multiresolutional level is represented by the coefficient that created the error. Figure 3 is a flow diagram of one embodiment of a process for performing error diffusion of a multiresolutional representation of an image in which an error filter is adapted to a multiresolutional level represented by a multiresolutional coefficient creating the error. In this embodiment, each multiresolutional coefficient is quantized and an error value(s) is passed to spatially neighboring multiresolutional coefficients as specified by an error filter adapted to a multiresolutional level represented by a multiresolutional coefficient creating the error, shown in process block 401.

[0056] The actual filter design used (e.g., in an image duplication device such as a copier) may, but need not, depend on the characteristics of the image duplication device, on some pre-filtering, or on the needs of an audience (e.g., sharper text or

smoother background may be required). A desired filter can be designed for the wavelet coefficients. The performance of a filter may, however, depend on the dots per inch, the device characteristics (including dot gain and isolated dot performance), and/or the intended purpose (e.g., text photocopying, image reproduction, etc.)

[0057] One of the keys to filter design for a filter to be used in the present invention is its extent. A filter that has a large spatial extent (and correspondingly small filter taps) will create sharper looking images because if there is an edge it is sure to express itself without being buried by the error of its immediate neighbor. However, this sharpness may also manifest itself as white or black shadow, or hemming. This white shadow that surrounds gray or black text both creates an artificial look and is a transient effect, but is also visually responsible for the appearance of enhanced sharpness. This effect of spatially larger filters holds true when doing error diffusion on wavelet coefficients.

[0058] However, larger spatial filters tend to also create a 'messier' or more random looking image. This effect is true for both pixel and wavelet domains, and in certain cases it counteracts the sharpening effect because thin lines may get "lost" in the messiness of a gray background that is not smoothly reproduced. An adaptive choice of filters can be made, for instance, large filters can be used in dark or bright region and small filters can be used in mid-tone regions. Extending the filter to the left of the pixel being quantized tends to result in less worm artifacts.

[0059] Spatial filtering may also be accomplished, according to one embodiment, by passing on the error differently to neighboring coefficients of different multiresolutional

levels. The taps of the spatial filter might then differ depending on the level. This is the same idea as above in Figure 3. For instance, each error might only be passed to close neighbors of the same multiresolutional level as the error-producer.

[0060] Figure 4 is a flow diagram of one embodiment of a process for performing error diffusion that includes passing the error generated using an error filter that depends on the subband of the coefficient whose error is being passed. Figure 5 is a flow diagram of one embodiment of a process for performing error diffusion that includes passing the error generated using an error filter that depends both on the subband and on the multiresolutional level of the coefficient whose error is being passed.

[0061] According to one embodiment, if the multiresolutional transform was overcomplete (e.g., taking $N \times N$ pixels to more than $N \times N$ transform coefficients) then the resulting halftone image would naturally include some interpolation (as each transform coefficient would result in a halftone coefficient). However, another way to handle this, according to an alternative embodiment, would be to only record uniformly sampled $N \times N$ of the halftone coefficients, that is, error diffuse over all the transform coefficients but only record a uniform sampling of $N \times N$ of the halftone coefficients (this is equivalent to downsampling the error diffused overcomplete representation).

[0062] When two close frequencies are superimposed, an unwanted lower frequency "beat" may result. This effect is called "moiré." This has been called "the most noticeable spatial domain problem in color printing." If the original image has a periodic structure (for example, if it has a regular cluster dot dither halftone structure

already in it) then error diffusion may lead to some moiré. However, when error diffusion is performed on the wavelet domain, there may be a regular periodic structure to the arrangement of the wavelet coefficients, and this greater periodicity may increase the risk of moiré. Non-regular arrangements of the transform coefficients should reduce the risk of moiré.

[0063] Figure 6A is a block diagram of one embodiment of an image processing device to perform error diffusion of a multiresolutional representation of an image. In one embodiment, the image processing device 601 is part of a printer. In an alternative embodiment, the image processing device 601 is part of a copier. Image processing device 601 comprises a receiving unit 602 to receive an image as a pixel representation. Image processing device 601 also comprises a processing unit 603 coupled with the receiving unit 601. Processing unit 602 converts the image from a pixel representation to a multiresolutional representation having multiresolutional coefficients. Processing unit 602 also quantizes each multiresolutional coefficient and passes error values in each iteration to spatially neighboring multiresolutional coefficients as specified by an error filter. Image processing device 601 further comprises an outputting unit 604 coupled with processing unit 603. Outputting unit 604 outputs a resulting quantized image.

[0064] Figure 6B is an alternative embodiment of an image processing device to perform error diffusion on a multiresolutional representation of an image. Referring to Figure 6B, image 610 is received by two-dimensional (2D) multiresolution transform 621. The 2D multiresolutional transform 621 converts image 610 to wavelet transform

coefficients. The coefficients are output to an optional processing unit 622. Processing 622 is optional and may perform other image processing such as, for example, but not limited to, JPEG 2000 compression or other compression, enhancement, and/or color management. These operations are performed on the coefficients.

[0065] Arranger 623 arranges the multiresolutional coefficients as described above to represent a spatial dependence of each multiresolutional coefficient. Error diffusion unit 624 receives the output from arranger 623 and performs the error diffusion described above on the arranged coefficients to produce a halftone image 625.

[0066] The method and apparatus disclosed herein may be integrated into advanced Internet- or network-based knowledge systems as related to information retrieval, information extraction, and question and answer systems. The system has a processor coupled to a bus. Also shown coupled to the bus is a memory which may contain instructions. Additional components coupled to the bus are a storage device (such as a hard drive, floppy drive, CD-ROM, DVD-ROM, etc.), an input device (such as a keyboard, mouse, light pen, bar code reader, scanner, microphone, joystick, etc.), and an output device (such as a printer, monitor, speakers, etc.). Of course, an exemplary computer system could have more components than these or a subset of the components listed.

[0067] The method described above can be stored in the memory of a computer system (e.g., set top box, video recorders, etc.) as a set of instructions to be executed. In addition, the instructions to perform the method described above could alternatively be stored on other forms of machine-readable media, including magnetic and optical disks.

For example, the method of the present invention could be stored on machine-readable media, such as magnetic disks or optical disks, which are accessible via a disk drive (or computer-readable medium drive). Further, the instructions can be downloaded into a computing device over a data network in a form of compiled and linked version.

[0068] Alternatively, the logic to perform the methods as discussed above, could be implemented in additional computer and/or machine readable media, such as discrete hardware components as large-scale integrated circuits (LSI's), application-specific integrated circuits (ASIC's), firmware such as electrically erasable programmable read-only memory (EEPROM's); and electrical, optical, acoustical and other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

[0069] Although the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense. Beyond the immediate practical possibilities for this technique, the idea of sigma-delta modulation on the wavelet domain is an interesting proposition which opens up new interpretations and may impact the general oversampled analog to digital conversion problem and audio applications.